

Median-of- k Jumplists

Markus E. Nebel* Elisabeth Neumann† Sebastian Wild†

September 29, 2016

We extend randomized jumplists introduced by Brönnimann et al. [1] to choose jump-pointer targets as median of a small sample, and present randomized algorithms with expected $O(\log n)$ time complexity that maintain the probability distribution of jump pointers upon insertions and deletions. We analyze the expected costs to search, insert and delete a random element. The resulting data structure, randomized median-of- k jumplists, is competitive to other dictionary implementations and supports particularly efficient iteration in sorted order and selection by rank. We further show that omitting jump pointers in small sublists hardly affects search costs, but significantly reduces the memory consumption. If space is tight and rank-select is needed, median-of- k jumplists are a promising option in practice.

1. Introduction

Jumplists were introduced by Brönnimann et al. [1] as a simple randomized dictionary implementation that allows to iterate over the stored elements in sorted order in a particularly efficient fashion. Furthermore, the search of any specific element is possible in logarithmic time. The core idea is to augment a sorted (singly-)linked list with additional shortcut pointers, the so-called *jump pointers*, to allow efficient access to all elements in the list. To obtain a binary-search-like behavior, a well-nested structure is required, i.e., jump-pointers may not cross each other. Figure 1 shows an exemplary jumplist.

Following the underlying linked-list pointers, iteration in sorted order is as efficient as it can be. Efficient search requires a favorable configuration of the jump pointers. However, we need a rule for choosing their targets that also allows to efficiently keep the jumplist in a favorable state when elements are inserted or deleted. Brönnimann et al. [1] proposed a randomized version, where jump pointers are set (from left to right) such that they invariably have a uniform distribution over all subsequent elements up to the first one that

*Technische Fakultät, Universität Bielefeld, Germany
nebel@techfak.uni-bielefeld.de

†Fachbereich Informatik, Technische Universität Kaiserslautern, Germany
{wild,neumann}@cs.uni-kl.de




Figure 1: A typical ordinary jumplist on $n = 30$ keys. Gray arrows indicate backbone (next) links, red thick arrows are jump pointers. The dotted green arrows indicate the end of a nodes sublist; these are not actually stored as pointers, but are implicitly given by the recursive structure. The used parameters (introduced later) are $k = 1$ and $w = 2$.

is already the target of another jump pointer. Accordingly, the outermost jump pointer, stored in the node of the smallest key, has the same probability to point to any element in the list, thereby dividing the list into two parts (called next- and jump-sublist). Since pointers may not cross, these parts are independent, and we require the same uniform distribution recursively in the two sublists.

In this note we generalize jumplists to use a more balanced distribution, where each jump pointer points to a median of a small sample of k elements of its sublist, instead of a single uniformly chosen one. (The original jumplists correspond to setting $k = 1$.) In both cases it is essential that insert and delete operations keep the probability distribution invariant, i.e., after a modification the distribution of jump pointers has to be the same as if all pointers were drawn from scratch. Building on the algorithms from Brönnimann et al. we present such insertion and deletion algorithms for median-of- k jumplists.

As for the original jumplists, inserting or deleting an element needs logarithmic time in expectation, including the cleanup to restore the distribution, making median-of- k jumplists competitive to other randomized dictionary implementations (see Section 1.1 for details). To quantify the influence of k , we analyze precisely the costs of searching, inserting and deleting elements. This allows to study the trade-off between a more efficient search resulting from better choices for the jump pointers and the resulting increased costs for the cleanup.

We furthermore introduce a novel search strategy (called spine search) that reduces the number of needed key comparisons. A further modification of jumplists discussed in this extended abstract is to omit jump pointers for sublists of small sizes (smaller than a threshold w). This allows to trade space for time: elements in these small sublists do not have to store a jump pointer, but the corresponding subfile can only be searched linearly.

The rest of this paper is structured as follows: In the following subsection we compare jumplists to similar data structures. In Section 2 and 3, we introduce ordinary jumplists. We present our spine search strategy in Section 4. Section 5 introduces the median-of- k extension, and Section 6 describes insertion and deletion for these. Our analysis results are given in Section 7, and we conclude the paper with a discussion of these result.

1.1. Related Work

Since binary search trees (BSTs) perform close to optimal on average—actually, even with high probability—if built from a random permutation [6], but poorly in the worst case, a natural idea is to enforce the average behavior through randomization. The most direct application of this paradigm is interestingly a relatively recent source: Martínez and Roura [8] devised efficient randomized insert and delete operations for BSTs, so that the distribution of the resulting trees is exactly the same as if the insertion order would have been a random permutation. This idea can even be made to work when equal keys are allowed in the tree [12].

Randomized BSTs have to store subtree sizes to maintaining the distribution, and so support rank-based operations without additional cost. This is not the case for the treaps of Seidel and Aragon [15], which store a random priority with each node instead. Treaps remain in random shape by enforcing a heap order w.r.t. the random priorities; apart from that their efficiency characteristics are very similar to randomized BSTs and hence, as we will see, to jumplists.

Unless further memory is used, BSTs do however not offer constant-time successor queries. Like jumplists, Pugh’s skip lists [13] are augmented sorted linked lists and thus successors can be found by simply following one pointer. Skip lists extend the list elements by towers of pointers of different heights, where each tower cell points to the next element in the list with a tower of at least this height. Using geometrically distributed tower heights, one achieves expected logarithmic search time with expected linear additional space. In practice, the varying tower heights can be inconvenient; jumplists achieve very similar performance without this complication.

An alternative to randomization is of course to enforce a certain minimal balance in BSTs [6]. Munro et al. [9] transfer the height-balancing effect of 2-3 trees to skip lists, and Elmasry [3] applied the weight-balancing criterion of $BB[\alpha]$ trees [11] to jumplists. Note that the latter achieves logarithmic update time only in an amortized sense.

2. Jumplists

We now present our (consolidated) definition of jumplists. It differs in a few details from the original version of Brönnimann et al. [1]; we discuss these differences in Appendix A.

Jumplists consist of *nodes*, where each node v stores a successor pointer ($v.next$) and a key ($v.key$). The nodes are connected using the next pointers to form a singly-linked list, the *backbone* of the jumplist, so that the key fields are sorted ascendingly.¹ For several reasons, it is convenient to add a “dummy” header node whose key field is ignored; if convenient we assume $v_0.key = -\infty$, but we never actually access this key in our algorithms.

If $x_1 < \dots < x_n$ are the keys stored in the jumplist, we thus have nodes v_0, v_1, \dots, v_n so that for $i = 1, \dots, n$ holds $v_i.key = x_i$ and $v_{i-1}.next = v_i$. A jumplist on n keys will always have $m = n + 1$ nodes, and we reserve these variables for this purpose throughout this paper: we write n if we count keys and m if we count nodes.

¹We consider jumplists as a representation of a *set* of keys; i.e., we assume that the keys stored in a jumplist are always distinct. Our operations will make sure that no duplicate insertions are possible.

Nodes in our jumplists actually come in two flavors: *plain nodes* only have the next and key fields, whereas *jump nodes* additionally store a *jump pointer* ($v.jump$) and an integer ($v.nsize$, discussed later). The jump pointer always points to an element further ahead in the list, and we require the following two conditions.

1. **Non-degeneracy:** Any node may be the target of at most one jump pointer, and jump pointers never point to the direct successor in the backbone.
2. **Well-nestedness:** Let $v \neq u$ be two jump nodes where v comes before u in the backbone, and let v^* resp. u^* be the nodes their jump pointers point to. (Note that $v^* \neq u^*$ by the first property). Then these nodes must appear in one of the following orders in the backbone: $u \dots v \dots v^* \dots u^*$ or $v \dots v^* \dots u \dots u^*$. (The second case allows $v^* = u$.)

Visually speaking, jump pointers may not cross.

The node types (plain or jump) are determined by the following rule that depends on a parameter $w \geq 2$, the *leaf size*: If the number of nodes in the jumplist is $m \leq w$, then all nodes are plain nodes. Otherwise the first node in the backbone, v_0 , is a jump node, and its jump pointer divides the rest of the jumplist into two parts: the nodes between v_0 and $v_0.jump$ (both exclusive) form the *next-sublist*, denoted by \mathcal{J}_1 , and the nodes from $v_0.jump$ up to the end of the list form the *jump-sublist* \mathcal{J}_2 . We now require the same node-type rule for these two sublists. The reader is invited to consult the example in Figure 1 to see this rule in action.

Sublists. This recursive rule is prototypical for many properties we will encounter later. To simplify their description, we associate with each node v the sublist whose head it is: The *sublist of v* starts at v (inclusive) and ends just before the first node targeted by a pointer originating before v , or continues to the end of the whole list if no such overarching jump pointer exists. As for the overall list, v acts as dummy header to its sublist: we do not count $v.key$ to the set of keys associated with v 's sublist. We also speak of the next- and jump-sublists of v , which are precisely the sublists of $v.next$ resp. $v.jump$. Figure 2 exemplifies these definitions. We include an imaginary “end pointer” in the figures, drawn as dotted green line, that connects a jump node with the last node in that node's sublist.

We will often be concerned with sublist sizes, so we use J_r for the number of nodes in \mathcal{J}_r , $r = 1, 2$. Indeed, our algorithms will also need access to the sublist sizes, therefore for every jump node v , we store in $v.nsize$ the number of nodes in the next-sublist of v .

Randomized Jumplists. The best pointer configuration (w.r.t. average search costs) would always cut the corresponding sublist in half; then searching the list would behave like a binary search in a sorted array. This best case is expensive to maintain upon insertions and deletions, therefore we resort to randomization (cf. [8, 1]): in a *randomized jumplist* the target of the topmost jump-pointer is drawn *uniformly* from all $m - 2$ possible targets, and conditional on the choice for the topmost pointer, we require the same property recursively and independently in both sublists. Formally, the probability $p(\mathcal{J})$




Figure 2: Illustration of the sublist definitions. The sublist of node 2 contains $m = 7$ nodes, but only the $n = m - 1 = 6$ keys $3, \dots, 8$ are considered to be stored in it. (The key of the header is always ignored within a sublist.) The sizes (number of nodes) of the next- and jump-sublist of node 2 are $J_1 = 2$ and $J_2 = 4$, respectively.

of a particular (legal) pointer configuration \mathcal{J} is

$$p(\mathcal{J}) = \begin{cases} 1, & m \leq w; \\ \frac{1}{m-2} \cdot p(\mathcal{J}_1) p(\mathcal{J}_2), & m > w. \end{cases} \quad (1)$$

3. Dangling-Min BSTs

The nature of jumplists—a backbone plus jump pointers—has technical advantages, e.g., for iteration, but it obfuscates an intimate relation between jumplists and binary search trees. The recursive structure of nested sublists obviously resembles a binary tree, but there is more to it than a superficial match of concepts. In this section, we define an (admittedly somewhat peculiar) variant of binary search trees (BSTs), the *dangling-min BSTs*, which turn out to be *isomorphic* to our jumplists (see below). The authors found reasoning about jumplists much more intuitive in terms of these trees, therefore we describe them here in some detail.

Let x_1, \dots, x_n be a sequence of n (distinct) keys from an ordered universe. The *dangling-min BSTs with leaf size $w \geq 2$* for this sequences is

- a leaf with the keys in sorted order if $n \leq w - 1$,
- or otherwise has as *root* a node containing *two* keys: the smallest key, $\min\{x_1, \dots, x_n\}$, as its *dangling minimum*, and the first key of the sequence after this smallest key has been removed as *root key* (i.e., the root key is x_1 , unless x_1 is the min; then we take x_2).

We now split the remaining sequence (without root key and min) into the subsequences of keys smaller resp. larger than the root key, preserving relative order, build dangling-min BSTs from both. These are attached as left resp. right subtree of the root.




Figure 3: The dangling-min BST with $w = 2$ for the sequence 11, 2, 5, 3, 1, 4, 10, 8, 7, 9, 6, 12, and the jumplist it corresponds to the tree.

Figure 3 shows an example of such a tree.

Next, we consider how to transform a jumplist to a dangling-min BST. Recall that a jumplist with m nodes actually stores only $n = m - 1$ keys, since the first node v_0 does not carry a key. If now $m \leq w$, v_0 is a plain node and the corresponding dangling-min BST is a leaf containing the $m - 1 \leq w - 1$ nodes in the jumplist. Otherwise, when $m > w$, v_0 is a jump node. Let x_1 be the key $v_1.next$ is x_j be the key in $v_0.jump$. Then, the root of the dangling-min BST has root key x_j and dangling min x_1 , and the left and right subtrees are generated recursively from the next- resp. jump-sublist of *head*.

Procedures like this can be concisely specified in an intuitive graphic syntax that we will use for our more intricate algorithms later. To illustrate the concept, the above transformation would be written as follows.

$$\begin{aligned}
 \text{MINBST} \left(\boxed{} \rightarrow \boxed{x_1 \dots x_n} \right) &= \boxed{x_1, \dots, x_n} & (m \leq w) \\
 \text{MINBST} \left(\boxed{} \rightarrow \underbrace{\boxed{x_1} \boxed{}}_{\mathcal{J}_1} \rightarrow \underbrace{\boxed{x_j} \boxed{}}_{\mathcal{J}_2} \right) &= \begin{array}{c} \boxed{x_1} \quad \boxed{x_j} \\ \swarrow \quad \searrow \\ \text{MINBST}(\mathcal{J}_1) \quad \text{MINBST}(\mathcal{J}_2) \end{array} & (m > w)
 \end{aligned}$$

The first equation defines MINBST on small jumplists ($m \leq w$); it shows a header without jump pointer, i.e., a plain node. The second equation defines MINBST on larger jumplists. Whenever variables appear on the left side, this is meant as a pattern to be matched to the actual input. The parts that match the variables are then used on the right-hand side.

Figure 3 also shows the jumplist corresponding to the given tree. As a larger example for the bijection between jumplists and dangling-min BSTs, Figure 4 shows the tree corresponding to the jumplist from Figure 1. One can show that a dangling-min BST built from a random permutation of $\{1, \dots, n\}$ has the same distribution as the dangling-min BST corresponding to a randomized jumplist.

4. Spine Search

As already sketched in Section 2, searching a jumplist is straightforward: if we are at node v and the key stored in $v.jump$ is not too large, we use the jump-pointer. Otherwise, the

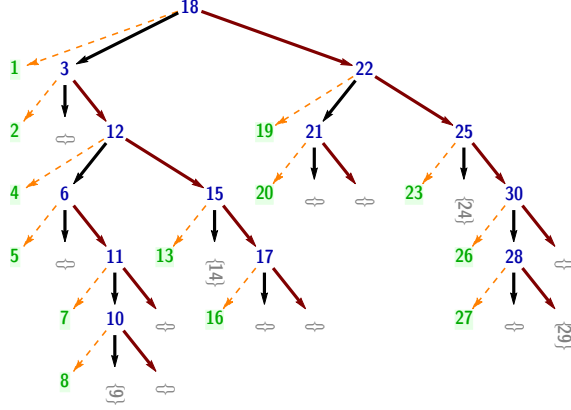


Figure 4: The dangling-min BST for the jumplist from Figure 1. Black arrows are left child pointers, red arrows are right child pointers, and the dotted yellow connect the subtree minimum to the root of a subtree. Gray nodes are leaves that contain between zero and $w - 1 = 1$ keys.

next-pointer is followed and the element stored in $v.next$ is compared to the searched key x . We call this strategy the classic search in the sequel. Note that there is a symmetric alternative: first compare x to $v.next$ and then with $v.jump$ (if needed). Brönnimann et al. [1] studied both and found that the second one needs more comparisons on average.

However, there is yet another search strategy considered neither in [1] nor [3] that is apparent in the dangling-min-BST representation. Consider searching key 8 in the list resp. the equivalent tree given in Figures 1 resp. 4. A classic search in this list inspects keys 18, 1, 3, 12, 4, 6, 11, 7, 10, 8 in the given order; a total of 10 key comparisons. Every step in the search that follows the next-pointer needs two comparisons—a general fact not only valid in our example.

Now do the search for 8 in the dangling-min BST as if it would be a regular BST (ignoring the subtree minima and stopping at the leaves). While doing so, we compare with keys 18, 3, 12, 6, 11, 10. All these steps need only one key comparison even though mostly the same keys are visited as above. However, our search is not yet finished; the reached leaf contains only 9, and we would (erroneously!) announce that 8 is not in the dictionary. Instead we have to return to the *last node we entered through right-child pointer* and inspect all the dangling mins along the “left spine” of the corresponding subtree. In our example, we return to 11 and make comparisons with 7 and 8, terminating successfully. We call this search strategy *spine search*. In our example, it needed 2 comparisons less than the classic search.

Spine search only compares x with the dangling-mins for nodes on the left spine, whereas the classic strategy does so for *any* node we leave through the left-child edge. Our modification is justified because when proceeding at node v to the right child we know that all keys left to v are smaller than x and thus x cannot be any of the dangling minima we skipped so far without comparison. Algorithm 1 gives our implementation of spine search for jumplists; after the BST-style search, we go back to the last safe checkpoint and continue a linear search from there.

Algorithm 1. Spine search in jumplists

```

SPINESearch(head, x)
  // Returns last node with key < x.
  // Assumes a sentinel node with key  $+\infty$  at end of backbone.
1  lastJumpedTo := head
2  repeat // BST-style search
3    if head.jump.key < x
4      head := head.jump; lastJumpedTo := head
5    else
6      head := head.next
7    end if
8  until head is PlainNode
9  head := lastJumpedTo
10 while head.next.key < x do head := head.next end while
11 return head

```

The left spine along which we compare with dangling mins is always a subset of the points where we take a left child edge on our search, so spine search never needs more comparisons than the classic strategy. On average, we expect that SPINESearch needs roughly as many key comparisons as the search in an ordinary BST over the same keys, since for most keys, the left spine is very short. Indeed, we prove in Appendix C that the linear search along the left spine only contributes to lower order terms when averaging over all possible unsuccessful searches—SPINESearch needs $2\ln(n)$ comparisons in the asymptotic average, compared to $3\ln(n)$ comparisons for the classic jumplist search strategy.

Since it is well known that search costs for BSTs benefit from locally balancing the trees by choosing medians as the subtrees' roots [4, 7, 2], it makes obviously sense to try the same optimization for jumplists. This is what we do next.

5. Median-of- k Jumplists

Let $k = 2t + 1$ for $t \in \mathbb{N}_0$ be a fixed parameter and assume $w \geq k + 1$. A randomized median-of- k jumplist is similar to an ordinary randomized jumplist as introduced in Section 2, but now we assume that the jump target is chosen as the median of k elements from the sampling range. More precisely, we define the probability of a (legal) jump pointer configuration \mathcal{J} as

$$p(\mathcal{J}) = \begin{cases} 1 & m \leq w; \\ \frac{\binom{J_1-1}{t} \binom{J_2-1}{t}}{\binom{m-2}{k}} \cdot p(\mathcal{J}_1) p(\mathcal{J}_2), & m > w. \end{cases} \quad (2)$$

This puts more probability on more balanced pointer configurations, and hence improves the expected search costs. Figure 5 shows a typical median-of-3 jumplist. A possible generalization could use asymmetric sampling with $\mathbf{t} = (t_1, t_2)$ and $k = t_1 + t_2 + 1$, where





Figure 5: A typical median-of-three ($k = 3$, $w = 4$) jumplist on $n = 30$ keys and its corresponding fringe-balanced dangling-min BST.

we select the $(t_1 + 1)$ -smallest instead of the median. Then, we have $\binom{J_1-1}{t_1}$ and $\binom{J_2-1}{t_2}$ in Equation (2). For the present work, we will however stick to the case $t_1 = t_2 = t$.

The situation is illustrated in the following sketch for $k = 3$ and $m = 10$; to have x_6 as the median of three elements from the sample range, we must select $t = 1$ further elements from $\{x_2, \dots, x_5\}$ and $t = 1$ further elements from $\{x_7, \dots, x_9\}$.



The number of such samples is precisely $\binom{J_1-1}{t} \binom{J_2-1}{t}$, which we have to divide by the total number of possible samples, $\binom{m-2}{k}$.

There is an alternative description of the distribution of the subproblem sizes, that is more convenient for our analysis. If we denote by $I_1 = J_1 - t - 1$, I_1 has a so-called beta-binomial distribution, $I_1 \stackrel{D}{=} \text{BetaBin}(m - 2 - k, t + 1, t + 1)$. (We write $\stackrel{D}{=}$ to denote equality in distribution.) The advantage lies in a stochastic representation: If $D_1 \stackrel{D}{=} \beta(t + 1, t + 1)$ (the beta distribution), we have conditional on D_1 that $I_1 \stackrel{D}{=} \text{Bin}(m - 2 - k, D_1)$; see also Appendix C.

Fringe-Balanced Dangling-Min BSTs. If we transform the distribution of median-of- k jumplists to dangling-min BSTs, we obtain *randomized fringe-balanced dangling-min BSTs*. The tree for a given permutation of n keys is now defined as follows: If $n \leq w - 1$, the tree is a leaf containing the sorted keys. Otherwise, remove the minimum from the list and make it the dangling min of the root. The root key is the median of the first k elements in the remaining sequence; it is also removed from the sequence. The subtrees are now built from the elements smaller resp. larger than the root in the remaining sequence. Again, we obtain the same distribution if we built the tree from a random permutation or a random median-of- k jumplist.

6. Insert and Delete Algorithms

In the following sections, we describe our insert and delete algorithms in prose and in the above introduced graphic syntax. We use the following conventions:

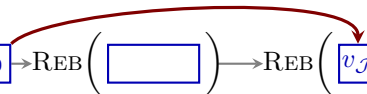
- The input of the algorithms, the “old” jumplist, is drawn as rectangle (or abbreviated by \mathcal{J}).
- A sequence of (output) leaf nodes is depicted by a rectangle with rounded corners.
- The position of insertion resp. deletion is marked in red.
- If the algorithm makes a random choice, each outcome is multiplied with its probability, and all outcomes are added up.

6.1. Rebalance

Algorithm REBALANCE is used if a jumplist needs to be (re)built from scratch. It only uses the backbone of the argument, any existing jump pointers are ignored. In the base case, i.e., if the argument contains $m \leq w$ nodes, a linked list of plain nodes with the same keys is returned.

$$\text{REB}\left(\boxed{v_0 \dots v_n}\right) = \boxed{v_0 \dots v_n} \quad (m \leq w)$$

If \mathcal{J} contains $m > w$ nodes, v_0 must become a jump node and we have to draw a jump target from the sample range. Conceptually, a sample of k nodes is drawn and the median w.r.t. the keys is chosen. The same distribution can actually be achieved without explicitly drawing samples using a random variable $J \stackrel{\text{d}}{=} \text{BetaBin}(m - 2 - k, t + 1, t + 1) + t + 2$ (see Section 5). Then the node v_J is the jump target. After the jump-pointer of v_0 has been initialized, the resulting next- and jump-sublist are rebalanced recursively.

$$\text{REB}\left(\boxed{v_0 \dots v_n}\right) = \boxed{v_0} \rightarrow \text{REB}\left(\boxed{}\right) \rightarrow \text{REB}\left(\boxed{v_{\mathcal{J}} }\right) \quad (m > w)$$


6.2. Insert

INSERT in jumplists consists of three phases found in many tree-based dictionaries: (unsuccessful) search, insertion, and cleanup. Unless x is already present, the search ends at the node with the largest key (strictly) smaller than x . There we insert a new node with key x into the backbone.

The new node however does not have a jump-pointer yet. Furthermore, the new node might need to be considered as potential jump target of its predecessors in the backbone. Thus, for all the nodes that have the new node in their sublist, we need to restore the pointer distribution. This is carried out by RESTOREAFTERINSERT.

Let m be the number of nodes after the insertion, i.e., including the new node. If $m \leq w$, the new node remains a plain node within a list of plain nodes, and no cleanup is necessary. If $m = w + 1$ due to the insertion, v_0 , which was a plain node before, now

has to become a jump node. In this case, REBALANCE is called on \mathcal{J} and the insertion terminates.

$$\text{RESTINS}\left(\boxed{} \boxed{} \boxed{}\right) = \text{REB}\left(\boxed{} \boxed{} \boxed{}\right) \quad (m \leq w)$$

If $m > w + 1$, we first restore the pointer distribution of v_0 . Due to the insertion of a new node, the sample range now contains an additional node u . Note that u is not necessarily the newly inserted node; if the new key is the first or second smallest in \mathcal{J} , u is the former second node of \mathcal{J} .

If we, conceptually, wanted to draw pointers for \mathcal{J} anew, there are two possibilities: either u is part of the sample, namely with probability $p = \frac{k}{m-2}$, or u is not part of it.² In the first case, we have to rebalance all of \mathcal{J} . Conditional on the event that u is not in the sample, the existing jump pointer of v_0 has the correct distribution: it has been chosen as the median of a random sample not containing u .

In the algorithm, we thus rebalance \mathcal{J} with probability p , where we draw the jump pointer of v_0 conditional on u being part of the sample. Otherwise, v_0 's jump pointer can be kept, and we continue recursively in the uniquely determined sublist that contains the inserted node—that is, unless v_0 does not have a jump pointer yet, since v_0 is the newly inserted node. In that case, we simply *steal* the jump-pointer of its direct successor, v_1 , which has the correct conditional distribution. Now v_1 does not have a jump pointer, and we treat this case recursively, as if v_1 was the newly inserted node.

$$\begin{aligned} \text{RESTINS}(\mathcal{J}) &= p \cdot \boxed{v_0} \rightarrow \text{REB}\left(\boxed{}\right) \rightarrow \text{REB}\left(\boxed{v_j} \boxed{} \boxed{v_n}\right) \\ &+ (1-p) \cdot \left\{ \begin{array}{ll} \boxed{x} \rightarrow \text{RESTINS}\left(\boxed{} \boxed{} \boxed{}\right) \rightarrow \boxed{v_j} \boxed{} & \mathcal{J} = \boxed{x} \rightarrow \boxed{v_0} \rightarrow \boxed{} \rightarrow \boxed{v_j} \boxed{} \\ \boxed{v_0} \rightarrow \text{RESTINS}\left(\boxed{} \boxed{} \boxed{}\right) \rightarrow \boxed{v_j} \boxed{} & \mathcal{J} = \boxed{v_0} \rightarrow \boxed{} \rightarrow \boxed{} \rightarrow \boxed{v_j} \boxed{} \\ \boxed{v_0} \rightarrow \boxed{} \rightarrow \text{RESTINS}\left(\boxed{v_j} \boxed{} \boxed{}\right) & \mathcal{J} = \boxed{v_0} \rightarrow \boxed{} \rightarrow \boxed{} \rightarrow \boxed{v_j} \boxed{} \end{array} \right. \end{aligned}$$

6.3. Delete

DELETE has the same three phases as INSERT: first a (successful) search finds the node to be deleted, then we actually remove it from the backbone. Finally, RESTOREAFTERDELETION

²The overall number of k -samples from n items is $\binom{n}{k}$; if we forbid, say, item n , we have only $\binom{n-1}{k}$ choices left.

$$\binom{n-1}{k} / \binom{n}{k} = \frac{(n-1)!k!(n-k)!}{k!(n-k-1)!n!} = \frac{n-k}{n} = 1 - \frac{k}{n} \quad (3)$$

performs the cleanup: the pointer distribution of those nodes whose sublists contained the deleted node has to be restored since their sample range has shrunk.

Let m be the number of nodes after deletion, and let u be the deleted node. We first assume that $u \neq v_0$; the case of deleting v_0 will be addressed later. If $m \leq w - 1$, \mathcal{J} is a list of plain nodes and can remain unaltered. If $m = w$, the size dropped from $w + 1$ to w due to the deletion, so v_0 has to be made a plain node.

$$\text{RESTDEL}\left(\boxed{\text{red box}}\right) = \boxed{v_0 \dots v_n} \quad (m \leq w)$$

Otherwise ($m > w$), v_0 is a jump node whose sublist contained u . There are two possible cases: either the sample drawn to choose $v_0.\text{jump}$ contained u , or not. In the latter case, the deletion of u does not affect the choice for $v_0.\text{jump}$ at all, and we recursively cleanup the uniquely determined sublist that formerly contained u . If u was indeed part of the sample, we have to rebalance \mathcal{J} .

It remains to determine the probability p that u was in the sample that led to the choice of $v_0.\text{jump}$. Unlike for insertion, p now depends on these two nodes. Let J_1 resp. J_2 be the sizes of the next- resp. jump-sublist *before* deletion; recall that we store J_1 in $v_0.\text{nsiz}$. Then p is given by the following expression:

$$p = \begin{cases} 1, & \text{if } u = v_0.\text{jump}; \\ \frac{t}{J_1 - 1}, & \text{if } u \text{ was in next-sublist (where } \frac{0}{0} := 1 \text{ in case } t = J_1 - 1 = 0); \\ \frac{t}{J_2 - 1}, & \text{if } u \text{ was in jump-sublist.} \end{cases} \quad (4)$$

In the case that the deleted node $u = v_0$, v_1 has become the new header, but its jump pointer now has the wrong distribution since v_0 's jump pointer no longer delimits its sample range. But observe that v_0 's (old) sample range was exactly v_1 's new sample range plus v_2 . Accordingly we only have to rebalance in case v_2 was part of the sample to select $v_0.\text{jump}$, which happens with probability $p = \frac{t}{J_1 - 1}$. Otherwise, we can conceptually impose v_0 's jump pointer on v_1 , which is easily implemented by swapping their keys, and continue the cleanup recursively in the next-sublist, as if v_1 had been deleted.

To conclude, the algorithm in both cases either rebalances the current sublist with probability p (as given above) and terminates, or it reuses the topmost old jump pointer and continues recursively.

$$\begin{aligned}
\text{RESTDEL}(\mathcal{J}) &= p \cdot \boxed{v_0} \rightarrow \text{REB} \left(\boxed{} \right) \rightarrow \text{REB} \left(\boxed{v_J} \boxed{v_n} \right) \\
&+ (1-p) \cdot \left\{ \begin{array}{ll} \boxed{v_1} \rightarrow \text{RESTDEL} \left(\boxed{\color{red}{}} \boxed{} \right) \rightarrow \boxed{v_j} \boxed{} & \mathcal{J} = \boxed{\color{red}{v_0}} \rightarrow \boxed{v_1} \rightarrow \boxed{} \rightarrow \boxed{v_j} \boxed{} \\ \boxed{v_0} \rightarrow \text{RESTDEL} \left(\boxed{} \boxed{\color{red}{}} \right) \rightarrow \boxed{v_j} \boxed{} & \mathcal{J} = \boxed{v_0} \rightarrow \boxed{} \boxed{\color{red}{}} \rightarrow \boxed{v_j} \boxed{} \\ \boxed{v_0} \rightarrow \boxed{} \rightarrow \text{RESTDEL} \left(\boxed{v_j} \boxed{\color{red}{}} \boxed{} \right) & \mathcal{J} = \boxed{v_0} \rightarrow \boxed{} \rightarrow \boxed{v_j} \boxed{\color{red}{}} \boxed{} \end{array} \right.
\end{aligned}$$

7. Analysis

We analyzed the expected behavior of median-of- k jumplists with leaf size w ; due to space constraints, we only list the results here. For `RESTOREAFTERINSERT` and `RESTOREAFTERDELETE`, the dominating costs is the call to `REBALANCE`, so we determine the expected size of the rebalanced sublist. Details and proofs are found in Appendix C.

Theorem 7.1:

Let a randomized median-of- k jumplist with leaf size w be given, where k and w are fixed constants. Abbreviate by $H(t) = H_{k+1} - H_{t+1}$ for H_n the harmonic numbers.

1. The expected number of key comparisons of a **spine search** is asymptotic to $1/H(t) \cdot \ln n$, as $n \rightarrow \infty$, when each position is equally likely to be requested.
2. The expected number of rebalanced elements in the **cleanup after insertion** is asymptotic to $k/H(t) \cdot \ln n$, as $n \rightarrow \infty$, when each of the $n+1$ possible gaps are equally likely.
3. The expected number of rebalanced elements in the **cleanup after deletion** is asymptotic to $k/H(t) \cdot \ln n$, as $n \rightarrow \infty$, when each key is equally likely to be deleted.
4. The expected number of additional machine words per key required to store the jumplist is asymptotic to $1 + \frac{2}{(w+1)H(t)}$ as $n \rightarrow \infty$.

8. Discussion

Our analysis shows that a search on average profits from sampling; in particular going from $k = 1$ to $k = 3$ entails significant savings: $\frac{12}{7} \ln n \approx 1.714 \ln n$ instead of $2 \ln n$ comparisons. As for median-of- k Quicksort, we see diminishing returns for much larger k . For jumplists, also the cleanup after insertions and deletions gets more expensive; it grows linear in k , so very large k will actually be harmful.

The efficiency of insertion resp. deletion depends on both, the time for the search and the time for the cleanup, so it is natural to ask for optimal k . Since the cost units




Figure 6: The k that minimizes the leading-term coefficient of total costs of insertion/deletion, if one comparison costs $\xi \in [0, 1]$ and each rebalanced element costs $1 - \xi$, i.e., $\arg \min_k \xi \cdot \frac{1}{H(t)} + (1 - \xi) \cdot \frac{k}{H(t)}$ as a function of ξ .

are rather different (comparisons vs. rebalanced elements) we need a weighing factor. Depending on the relative weight $\xi \in [0, 1]$ of comparisons, we can compute optimal k , see Figure 6. In the realistic range, we should try $k = 1, 3$, or 5 , unless we do many more searches than updates.

We also conducted a small running time study based on a Java implementation [17] of our new data structure that essentially confirms our analytical findings: Sampling indeed leads to small savings for searches, but slows down insertions and deletions even more. Comparing running times with that of Java’s TreeMap (a red-black tree implementation) proves our data structure competitive: for iterating over all elements jumplists are about 50% faster, searches are between 20% and 100% slower (depending on the choice for w) and for insertions/deletions TreeMaps are 5 to 10 times faster. However, we compared a proof of concept implementation to a highly engineered library data structure; further improvements might be attainable.

Most notably, however, TreeMaps use 4 additional words per key (without even storing subtree sizes), whereas our jumplists never need more than ~ 2.3 additional words per key. Indeed with $w \approx 100$, we use less than 1.04 additional words per key. For $n = 10^6$ keys, $w \approx 100$ did not affect searches much (+25%) but actually sped up insertions and deletions (roughly by a factor of 2!).

9. Conclusion

In this note we presented median-of- k jumplists and analyzed their average efficiency based on theoretical cost measures like comparisons and rebalanced elements. Jumplists are competitive to other dictionary implementations, and might be unique in combining very low memory footprint with efficient rank select and worst-case constant-time successor operations.

Some interesting questions are left open: What is the optimal choice for w ? Answering this question requires second-order terms of search, insertion and deletion costs; due

to the underlying mathematical challenges it is unlikely that those can be computed exactly, but an upper bound using analysis results on Quicksort should be possible. Other future directions might be the analysis of branch misses, in particular in the context of an asymmetric sampling strategy, and the design of a “bulk insert” algorithm that is faster than inserting elements subsequently, one at a time.

On modern computers the cache performance of data structures is important for their running time efficiency. Here, a larger fanout of nodes is beneficial since it reduces the expected number of I/Os. For jumplists this can be achieved by using more than one jump pointer in each node. The case of two jump pointers per node has been worked out in detail [10], but the general scheme invites further investigation.

References

- [1] H. Brönnimann, F. Cazals, and M. Durand. Randomized jumplists: A jump-and-walk dictionary data structure. In *STACS 2003*, pages 283–294, 2003. doi: 10.1007/3-540-36494-3_26.
- [2] M. Drmota. *Random Trees*. Springer, 2009. ISBN 978-3-211-75355-2.
- [3] A. Elmasry. Deterministic jumplists. *Nordic Journal of Computing*, 12(1):27–39, March 2005.
- [4] D. H. Greene. *Labelled formal languages and their uses*. Ph.D. thesis, Stanford University, January 1983.
- [5] P. Hennequin. Combinatorial analysis of Quicksort algorithm. *Informatique théorique et applications*, 23(3):317–333, 1989.
- [6] D. E. Knuth. *The Art Of Computer Programming: Searching and Sorting*. Addison Wesley, 2nd edition, 1998. ISBN 978-0-20-189685-5.
- [7] H. M. Mahmoud. *Pólya Urn Models*. Chapman & Hall, 2008. ISBN 978-1-4200-5983-0.
- [8] C. Martínez and S. Roura. Randomized binary search trees. *J. ACM*, 45(2):288–323, March 1998. doi: 10.1145/274787.274812.
- [9] J. I. Munro, T. Papadakis, and R. Sedgewick. Deterministic skip lists. In *ACM-SIAM Symposium on Discrete Algorithms*, SODA 1992, pages 367–375. SIAM, 1992. ISBN 0-89791-466-X.
- [10] E. Neumann. *Randomized Jumplists With Several Jump Pointers*. Bachelor’s thesis, 2015. URL <http://nbn-resolving.de/urn/resolver.pl?urn:nbn:de:hbz:386-kluedo-41642>.
- [11] J. Nievergelt and E. M. Reingold. Binary search trees of bounded balance. *SIAM Journal on Computing*, 2(1):33–43, 1973. doi: 10.1137/0202005.

- [12] T. A. Pasanen. Random binary search tree with equal elements. *Theoretical Computer Science*, 411(43):3867–3872, October 2010. ISSN 03043975. doi: 10.1016/j.tcs.2010.06.023.
- [13] W. Pugh. Skip lists: A probabilistic alternative to balanced trees. *Communications of the ACM*, 33(6):668–676, June 1990. doi: 10.1145/78973.78977.
- [14] S. Roura. Improved Master Theorems for Divide-and-Conquer Recurrences. *Journal of the ACM*, 48(2):170–205, 2001.
- [15] R. Seidel and C. R. Aragon. Randomized search trees. *Algorithmica*, 16(4-5):464–497, October 1996. doi: 10.1007/BF01940876. URL <http://link.springer.com/10.1007/BF01940876>.
- [16] S. Wild. *Dual-Pivot Quicksort and Beyond: Analysis of Multiway Partitioning and Its Practical Potential*. Doktorarbeit (Ph.D. thesis), Technische Universität Kaiserslautern, 2016.
- [17] S. Wild. sebwild/jumplist: snapshot-for-paper. September 2016. doi: 10.5281/zenodo.155326.

Appendix

A. Our Jumplists vs. the Original Version

Our definition of jumplists differs in some details from the original version. We list the differences here, and also discuss we decided to make these changes.

Symmetry. In the original version of the jumplist, the jump pointer is allowed to target any node from the sublist, except the header itself. Thus there are $m - 1$ possible choices. In this setting, the size of the next-sublist can attain any value between 0 and $m - 2$, whereas the size of the jump-sublist is between 1 and $m - 1$.

This asymmetry between the subproblems felt “wrong” to us, and is in fact easily avoided: We simply disallow the direct successor of the head as possible target. This modification restores symmetry between next- and jump-sublist: both must be non-empty and contain at most $m - 2$ nodes and their sizes have the same distribution. Moreover, forbidding the direct successor as jump target is also a natural requirement since such a degenerate “shortcut” is totally useless in searches.

Small Sublists. The original jumplists only have one type of nodes which corresponds to our jump node. In the case $m = 1$, Brönnimann et al. resort to assigning an “exceptional pointer” to the direct successor; note that this node actually lies outside (one behind) of the current sublist. These pointers are of no use, as they are never followed during (jump-and-walk) search.

We feel that in a typical implementation with heap-allocated memory for each node, different node types (and sizes) are not a problem; rather indeed a chance to *save* memory. We thus introduced the plain node without jump pointer, used whenever the sublist has at most w nodes. $w \geq 2$ is required if we want to avoid useless jump pointers that point to the direct successor.

This also allows to declare that every node has at most one incoming jump pointer, which we again see as a natural requirement, since from the perspective of a search starting at the header, such shortcuts with the same target are redundant (and will indeed never be taken).

Finally, tuning of parameter w allows us to adapt our data structure to the problem at hand. Storage space can be traded for search time.

Sentinel vs. Circularly closed. The original jumplist implementation has a circularly closed backbone, i.e., the next pointer of the last node in the list points to the overall header again. This is the classical pattern for terminating an (unsuccessful) search, that avoids special treatment of a list without stored keys.

Since the backbone is always sorted, we borrow from sorting algorithms another classic technique that is slightly more efficient here. We add a *sentinel* node with key $+\infty$ to the end of the list, so that we can omit any explicit boundary checks during searches.

B. Pseudocode

We give detailed pseudocode for the basic operations on median-of- k jumplists with leaf size w here. We first list the four procedures CONTAINS, INSERT, DELETE and RANKSELECT that constitute the public interface of the data structure; the other procedures can be thought of as low-level procedures typically hidden from the user of the data structure.

We assume that jumplists are represented using the following records/objects.

Used Objects/Structs

- 1 `JumpList(head, size)`
- 2 `PlainNode(key, next)`
- 3 `JumpNode(key, next, jump, nsize)`

References/pointers to nodes can refer to a PlainNode or to a JumpNode, and we assume there is an efficient method to check which type a particular instance has. If *node* is a reference to a PlainNode, we write *node.key* and *node.next* for the key-value and next-pointer fields of the referenced PlainNode; similarly for the other types.

CONTAINS(*jumpList*, *x*)

- ```
// Returns whether x is present in jumpList and how many elements $< x$ it stores.
1 (node, r) := SPINESEARCH(jumpList.head, x)
2 candidate := node.next
3 return (candidate.key == x, r)
```

INSERT(*jumpList*, *x*)

```

 // Insert x into jumpList; does nothing if x is already present.
1 (node, r) := SPINESEARCH(jumpList.head, x)
2 if node.next.key ≠ x // x not yet present
3 node.next := new PlainNode(x, node.next) // Add new node in backbone.
4 n := jumpList.size + 1; jumpList.size := n
5 jumpList.head := RESTOREAFTERINSERT(jumpList.head, n + 1, r + 1)
6 end if

```

DELETE(*jumpList*, *x*)

```

 // Removes x from jumpList; does nothing if x is not present.
1 (node, r) := SPINESEARCH(jumpList.head, x)
2 if node.next.key == x // x is present
3 delNode := node.next
4 node.next := delNode.next // Remove delNode from backbone.
5 n := jumpList.size - 1; jumpList.size := n
6 jumpList.head := RESTOREAFTERDELETE(jumpList.head, n + 1, r + 1, delNode)
7 end if

```

RANKSELECT(*jumpList*, *rank*)

```

 // Returns the element with (zero-based) rank rank, i.e., the (rank + 1)st smallest element.
1 head := jumpList.head; r := rank + 1
2 repeat
3 if r > head.nsize
4 r := r - (head.nsize + 1); head := head.jump
5 else
6 r := r - 1; head := head.next
7 end if
8 if r == 0 then return head.key end if
9 until head is PlainNode
10 repeat
11 r := r - 1; head := head.next
12 until r == 0
13 return head

```

The above methods make use of the following internal procedures. We give again a spine search implementation, that is augmented to determine also the rank of the element. Using the rank makes the procedures to restore the distribution key after insertions or deletions a bit more convenient to state, and also avoids re-doing key comparisons there.

The given implementation of SPINESEARCH, CONTAINS, INSERT, and DELETE assume a sentinel node *tail* at the end of the linked list that has *tail.key* =  $+\infty$ , i.e., a value larger than any actual key value; we do however not count *tail* towards the *m* nodes of a jumplist since *tail* can be shared across all instances of jumplists. The sentinel may never be the target of any jump pointer. We could avoid the need for the sentinel at the

expense of a null-check of the next pointer, before comparing the successor's key (line 12 in SPINESEARCH, line 2 in CONTAINS, line 2 in INSERT, and line 2 in DELETE). Since using the sentinel is a bit more efficient and gives more readable code, we stick to this assumption.

SPINESEARCH(*head*, *x*)

```

 // Returns last node with key < x and its zero-based rank,
 // i.e., the number of nodes with key < x
1 rank := 0; steppedOver := 0; lastJumpedTo := head
2 repeat // BST-style search
3 if head.jump.key < x
4 rank := rank + head.nsize + 1 + steppedOver
5 head := head.jump
6 steppedOver := 0; lastJumpedTo := head
7 else
8 head := head.next; steppedOver := steppedOver + 1
9 end if
10 until head is PlainNode
11 head := lastJumpedTo
12 while head.next.key < x // Linear search from lastJumpedTo
13 head := head.next; rank := rank + 1
14 end while
15 return (head, rank)

```

REBALANCE(*head*, *m*)

```

 // Draws jump pointers in for m nodes starting with head (inclusive)
 // according to the randomized jumplist distribution.
 // Returns new first (possibly still head) and last node of the sublist.
1 if m ≤ w
2 Replace head and its m − 1 successors by m linked PlainNodes.
3 return (new head, new end)
4 else
5 S := random k-element subset of [2..m − 1]
6 jumpIndex := MEDIAN(S) // S(t1+1) in general
7 return SETJUMPANDREBALANCE(head, m, jumpIndex)
8 end if

```

SETJUMPANDREBALANCE(*head*, *m*, *j*)

```

 // Rebalances the sublist starting at head containing m nodes,
 // where we fix the topmost jump pointer to point to the element of rank j.
 // Returns new first (possibly still head) and last node of the sublist.
1 (nextStart, nextEnd) := REBALANCE(head.next, j − 1)
2 (jumpStart, jumpEnd) := REBALANCE(nextEnd.next, m − j)
3 nextEnd.next := jumpStart
4 return (new JumpNode(head.key, nextStart, jumpStart, j − 1), jumpEnd)

```

RESTOREAFTERINSERT( $head, m, r$ )

```

 // Restore distribution in sublist with header $head$ and of size m after an insertion at position r .
 // m is the number of nodes in the sublist, including the new element
 // Returns the new head of the sublist (possibly still $head$).
1 if $m \leq w + 1$ // Base case
2 if $m == w + 1$ // We need a new JumpNode, so rebalance.
3 $(head, end) := \text{REBALANCE}(head, m)$
4 end if
5 else
6 $newElementInSample := \text{COINFLIP}(\frac{k}{m-2})$ // true with probability $\frac{k}{m-2}$
7 if $newElementInSample$ // Rebalance conditional on new index being in sample.
8 $newIndex := \max\{r, 2\}$
9 $S := \text{random } (k - 1)\text{-element subset of } [2..m - 1] \setminus \{newIndex\}$
10 $jumpIndex := \text{MEDIAN}(newIndex \cup S)$ // $S_{(t_1+1)}$ in general
11 $(head, end) := \text{SETJUMPANDREBALANCE}(head, m, jumpIndex)$
12 else // topmost jump pointer can be kept
13 if $r == 0$ // new node is head of sublist, so steal successor's jump.
14 // Swap roles of the two nodes.
15 $succ := head.next$
16 Swap key and $next$ fields of $head$ and $succ$.
17 $head := succ$
18 end if
19 $J_1 := head.nsize$
20 if $r \leq J_1 + 1$ // New element is in next-sublist.
21 $J_1 := J_1 + 1$; $head.nsize := J_1$
22 $succ := \text{RESTOREAFTERINSERT}(head.next, J_1, \max\{0, r - 1\})$
23 $head.next := succ$
24 else // New element is in jump-sublist.
25 $J_2 := m - 1 - J_1$
26 $jumpHead := \text{RESTOREAFTERINSERT}(head.jump, J_2, r - 1 - J_1)$
27 if $jumpHead \neq head.jump$ // Have to reconnect backbone
28 $head.jump := jumpHead$
29 $(lastInNext, rr) := \text{SPINESEARCH}(head, jumpHead.key)$
30 $lastInNext.next := jumpHead$
31 end if
32 end if
33 end if
34 return $head$

```

```

RESTOREAFTERDELETE(head, m, r, delNode)
 // Restore distribution in sublist with header head and size m after a deletion at position r.
 // m is the number of nodes in the sublist, excluding the just deleted element delNode.
 // Returns the new head of the sublist (possibly still head).
1 if m ≤ w
2 if m == w // head is a JumpNode, must become PlainNode
3 head := new PlainNode(head.key, head.next)
4 end if
5 else
6 $J_1 := \begin{cases} \textit{delNode}.\textit{nsiz}, & \textbf{if } r == 0; \\ \textit{head}.\textit{nsiz}, & \textbf{else.} \end{cases}$
7 $p := \begin{cases} 1, & \textbf{if } r == J_1 + 1; \text{ // deleted node was jump target} \\ [J_1 == 1 \wedge r \leq 1], & \textbf{else if } k == 1; \text{ // special case to avoid } \frac{0}{0} \\ \frac{t_1}{J_1 - 1}, & \textbf{else if } r < J_1 + 1; \\ \frac{t_2}{m - 1 - J_1}, & \textbf{else.} \end{cases}$
8 deletedElementInSample := COINFLIP(p) // true with probability p
9 if deletedElementInSample // Rebalance sublist.
10 (head, end) := REBALANCE(head, m)
11 else // Topmost jump pointer can be kept.
12 if r == 0 // Impose deleted head's pointer onto successor.
13 // Swap roles of the two nodes.
14 Swap key and next fields of head and delNode.
15 Swap head and delNode.
16 end if
17 if r < J1 + 1 // Deletion in next-sublist.
18 J1 := J1 - 1; head.nsiz := J1
19 succ := RESTOREAFTERDELETE(head.next, J1, max{0, r - 1}, delNode)
20 head.next := succ
21 else // Deletion in jump-sublist.
22 J2 := m - 1 - J1
23 jumpHead := RESTOREAFTERDELETE(head.jump, J2, r - 1 - J1, delNode)
24 if jumpHead ≠ head.jump // Have to reconnect backbone
25 head.jump := jumpHead
26 (lastInNext, rr) := SPINESEARCH(head, jumpHead.key)
27 lastInNext.next := jumpHead
28 end if
29 end if
30 end if
31 return head

```

## C. Analysis of Median-of-k Jumplists

In this appendix, we give the details for the analysis of the performance of median-of- $k = 2t + 1$  jumplists with leaf size  $w \geq k + 1$ . Before we start, we collect some notation and basic facts used in the analysis.

### C.1. Preliminaries

For the reader's convenience, we collect here a few preliminaries used in the analysis; this material is introduced in more detail in [16] in the context of the analysis of Quicksort.

#### C.1.1. Distributions

**Beta Distribution.** The *beta distribution* has two parameters  $\alpha, \beta \in \mathbb{R}_{>0}$  and is written as  $\text{Beta}(\alpha, \beta)$ . If  $X \stackrel{D}{=} \text{Beta}(\alpha, \beta)$ , we have  $X \in (0, 1)$  and it has the density

$$f(x) = \frac{x^{\alpha-1}(1-x)^{\beta-1}}{B(\alpha, \beta)}, \quad (5)$$

where  $B(\alpha, \beta) = \Gamma(\alpha)\Gamma(\beta)/\Gamma(\alpha + \beta)$  is the beta function.

**Beta-Binomial Distribution.** The *beta-binomial distribution* is a discrete distribution with parameters  $n \in \mathbb{N}_0$  and  $\alpha, \beta \in \mathbb{R}_{>0}$ . It is written as  $\text{BetaBin}(n, \alpha, \beta)$ . If  $I \stackrel{D}{=} \text{BetaBin}(n, \alpha, \beta)$ , we have  $I \in [0..n]$  and

$$\mathbb{P}[I = i] = \binom{n}{i} \frac{B(\alpha + i, \beta + (n - i))}{B(\alpha, \beta)}, \quad i \in \mathbb{Z}. \quad (6)$$

(Recall that the binomial coefficients are zero unless  $i \in [0..n]$ .) An alternative representation of the weights is:

$$\binom{n}{i} \frac{B(\alpha + i, \beta + (n - i))}{B(\alpha, \beta)} = \binom{n}{i} \frac{\alpha^{\bar{i}} \beta^{\overline{n-i}}}{(\alpha + \beta)^{\bar{n}}}. \quad (7)$$

Note that for the special case  $\alpha = \beta = 1$ , we obtain  $\text{BetaBin}(n, 1, 1) \stackrel{D}{=} \mathcal{U}[0..n]$ : for  $i \in [0..n]$  then holds

$$\mathbb{P}[I = i] = \binom{n}{i} \frac{1^{\bar{i}} 1^{\overline{n-i}}}{2^{\bar{n}}} = \frac{n!}{i!(n-i)!} \frac{i!(n-i)!}{(n+1)!} = \frac{1}{n+1}. \quad (8)$$

**Stochastic Representation.** There is a second way to obtain beta-binomial distributed random variables: we first draw a random probability  $D \stackrel{D}{=} \text{Beta}(\alpha, \beta)$  according to a beta distribution, and then use this as parameter of a binomial distribution, i.e.,  $I \stackrel{D}{=} \text{Bin}(n; d)$  *conditional* on  $D = d$ . In other words, the beta-binomial distribution is a *mixed* binomial distribution, using a beta *mixer*  $D$  to determine the parameter of the binomial. This justifies its name.

**Dirichlet-Multinomial Distribution.** The multivariate generalization of the beta distribution is known as the Dirichlet distribution, and the multinomial distribution likewise generalized the binomial distribution; one can then define the *Dirichlet-multinomial distribution* as the mixed multinomial distribution, where the probabilities are given by a Dirichlet distribution. This generalization does not play a big role in this work, but we use it as a mere syntactic shortcut: we write

$$\mathbf{I} \stackrel{\mathcal{D}}{=} \text{DirMult}(n, \alpha, \beta) \quad (9)$$

to mean that

$$I_1 \stackrel{\mathcal{D}}{=} \text{BetaBin}(n, \alpha, \beta), \quad (10)$$

$$I_2 = n - I_1, \quad (11)$$

$$\mathbf{I} = (I_1, I_2). \quad (12)$$

### C.1.2. Some Expectations

**Lemma C.1:** Let  $X \stackrel{\mathcal{D}}{=} \text{Bin}(n, p)$  for  $n \in \mathbb{N}_0$  and  $p \in (0, 1]$ . Then we have with  $q = 1 - p$  that

$$\mathbb{E}[X^{-1}] = \mathbb{E}\left[\frac{1}{X+1}\right] = n^{-1} p^{-1} \cdot (1 - q^{n+1}), \quad (13)$$

$$\mathbb{E}[X^{-2}] = \mathbb{E}\left[\frac{1}{(X+1)(X+2)}\right] \leq n^{-2} p^{-2}. \quad (14)$$

**Proof:** For  $m \in \{1, 2\}$ , we compute

$$\mathbb{E}[X^{-m}] = \sum_{x=0}^n x^{-m} \cdot \binom{n}{x} p^x q^{n-x} \quad (15)$$

$$= n^{-m} p^{-m} \sum_{x=0}^n \binom{n+m}{x+m} p^{x+m} q^{(n+m)-(x+m)} \quad (16)$$

$$= n^{-m} p^{-m} \sum_{x=m}^{n+m} \binom{n+m}{x} p^x q^{(n+m)-x} \quad (17)$$

$$\stackrel{[\text{binom. thm.}]}{=} n^{-m} p^{-m} \left( \underbrace{(p+q)^{n+m}}_{=1} - \sum_{x=0}^{m-1} \binom{n+m}{x} p^x q^{(n+m)-x} \right) \quad (18)$$

For the first part of the claim, we set  $m = 1$  and find that the sum reduces to  $q^{n+1}$ ; for the second part of the claim, we use  $m = 2$  and note that the expression in the outer parentheses is at most 1.  $\square$

### C.1.3. Distributional Master Theorem

We cite here a handy theorem that allows to obtain asymptotic solutions to most of the recurrences we encounter in the analysis of jump lists. The formulation below appears as Theorem 2.76 in [16]; it is essentially a shortcut to using the more general continuous master theorem of Roura [14].

**Theorem C.2 (Distributional Master Theorem (DMT)):**

Let  $(C_n)_{n \in \mathbb{N}_0}$  be a family of random variables that satisfies the distributional recurrence

$$C_n \stackrel{\mathcal{D}}{=} T_n + \sum_{r=1}^s A_r^{(n)} \cdot C_{J_r^{(n)}}^{(r)}, \quad (n \geq n_0), \quad (19)$$

where  $(C_n^{(1)})_{n \in \mathbb{N}}, \dots, (C_n^{(s)})_{n \in \mathbb{N}}$  are independent copies of  $(C_n)_{n \in \mathbb{N}}$ , which are also independent of  $\mathbf{J}^{(n)} = (J_1^{(n)}, \dots, J_s^{(n)}) \in \{0, \dots, n-1\}^s$ ,  $\mathbf{A}^{(n)} = (A_1^{(n)}, \dots, A_s^{(n)}) \in \mathbb{R}_{\geq 0}^s$  and  $T_n$ . Define  $\mathbf{Z}^{(n)} = \mathbf{J}^{(n)}/n$  and assume that the components  $Z_r^{(n)}$  of  $\mathbf{Z}^{(n)}$  fulfill uniformly for  $z \in (0, 1)$

$$n \cdot \mathbb{P}[Z_r^{(n)} \in (z - \frac{1}{n}, z]] = f_{Z_r^*}(z) \pm O(n^{-\delta}), \quad (n \rightarrow \infty), \quad (20)$$

for a constant  $\delta > 0$  and a Hölder-continuous function  $f_{Z_r^*} : [0, 1] \rightarrow \mathbb{R}$ . Then  $f_{Z_r^*}$  is the density of a random variable  $Z_r^*$  and  $Z_r^{(n)} \xrightarrow{\mathcal{D}} Z_r^*$ .

Let further

$$\mathbb{E}[A_r^{(n)} \mid Z_r^{(n)} \in (z - \frac{1}{n}, z]] = a_r(z) \pm O(n^{-\delta}), \quad (n \rightarrow \infty), \quad (21)$$

for a function  $a_r : [0, 1] \rightarrow \mathbb{R}$  and require that  $f_{Z_r^*}(z) \cdot a_r(z)$  is also Hölder-continuous on  $[0, 1]$ . Moreover, assume  $\mathbb{E}[T_n] \sim Kn^\alpha \log^\beta(n)$ , as  $n \rightarrow \infty$ , for constants  $K \neq 0$ ,  $\alpha \geq 0$  and  $\beta > -1$ . Then, with  $H = 1 - \sum_{r=1}^s \mathbb{E}[(Z_r^*)^\alpha a_r(Z_r^*)]$ , we have the following cases.

1. If  $H > 0$ , then  $\mathbb{E}[C_n] \sim \frac{\mathbb{E}[T_n]}{H}$ .
2. If  $H = 0$ , then  $\mathbb{E}[C_n] \sim \frac{\mathbb{E}[T_n] \ln n}{\tilde{H}}$  with  $\tilde{H} = -(\beta + 1) \sum_{r=1}^s \mathbb{E}[(Z_r^*)^\alpha a_r(Z_r^*) \ln(Z_r^*)]$ .
3. If  $H < 0$ , then  $\mathbb{E}[C_n] = O(n^c)$  for the  $c \in \mathbb{R}$  with  $\sum_{r=1}^s \mathbb{E}[(Z_r^*)^c a_r(Z_r^*)] = 1$ .

The following technical lemma is used repeatedly to show that (20) is satisfied. Noting that the subproblem sizes,  $J_1$  resp.  $J_2$ , for jumplists have the same distribution as the sizes of recursive calls in median-of- $k$  Quicksort (up to a constant offset), we can reuse previous work by one of the authors on that.

**Lemma C.3 (Convergence in Density):** Let  $J^{(n)} \stackrel{\mathcal{D}}{=} \text{BetaBin}(n + d, \alpha, \beta) + c$  as  $n \rightarrow \infty$ , where  $c$  and  $d$  are two fixed integer constants. Then

$$n\mathbb{P}[J^{(n)}/n \in (z - \frac{1}{n}, z]] = f(z) \pm O(n^{-1}) \quad (22)$$

where  $f(z) = z^{\alpha-1}(1-z)^{\beta-1}/B(\alpha, \beta)$  is the density function of the beta distribution with parameters  $\alpha$  and  $\beta$ . We say that  $J^{(n)}/n$  converges in density to  $\text{Beta}(\alpha, \beta)$ .

**Proof:** See [16, Section 6.3]. □



## C.2. Search Costs (Spine Search)

Let  $P_n$  be the (random) total costs to search all numbers  $x \in \{0.5, 1.5, \dots, n + 0.5\}$  in  $\mathcal{J}_n$  using SPINESEARCH; to be specific, we count here the number of key comparisons. The corresponding quantity in BSTs is called external path length, and we will likewise do so for  $P_n$ . By dividing  $P_n/n$ , we obtain the average costs of one call to SPINESEARCH when all  $n + 1$  gaps are equally likely to be requested.  $P_n$  is random w.r.t. to the locations of the jump pointers in  $\mathcal{J}_n$ .

To set up a recurrence for  $P_n$ , the perspective of random dangling-min BSTs is most convenient, since SPINESEARCH exactly follows this implicit tree structure.

We describe recurrences here in terms of the distributions of families of random variables.

$$P_n \stackrel{\mathcal{D}}{=} (n + 1) + (S_n + L_n + 1) + S_{J_1} + P_{J_1} + P_{J_2}, \quad (n \geq w), \quad (23)$$

$$P_n \stackrel{\mathcal{D}}{=} \frac{(n + 1)(n + 2)}{2}, \quad (n < w), \quad (24)$$

$$\mathbf{J} = (J_1, J_2) = \mathbf{I} + \mathbf{t}, \quad (25)$$

$$\mathbf{I} = (I_1, I_2) \stackrel{\mathcal{D}}{=} \text{DirMult}(n - 1 - k; t + 1, t + 1), \quad (26)$$

$$S_n \stackrel{\mathcal{D}}{=} 1 + S_{J_1}, \quad (n \geq w), \quad (27)$$

$$S_n \stackrel{\mathcal{D}}{=} 0, \quad (n < w), \quad (28)$$

$$L_n \stackrel{\mathcal{D}}{=} L_{J_1}, \quad (n \geq w), \quad (29)$$

$$L_n \stackrel{\mathcal{D}}{=} n, \quad (n < w). \quad (30)$$

The terms  $P_{J_1}$  and  $P_{J_2}$  on the right-hand side denote members of independent copies of the family of random variables  $(P_n)_{n \in \mathbb{N}_0}$ , which are also independent of  $\mathbf{J} = \mathbf{J}^{(n)}$ . (We omitted the superscripts above for readability.)

Two additional quantities of independent interest appear in this recurrence:  $S_n$  is the number of internal nodes on the “left spine” of the tree; this is also the depth of the internal node with the smallest key (excluding dangling mins, of course; the smallest key overall is of course always in the root). For ordinary BSTs,  $S_n$  is essentially the number of left-to-right minima in a random permutation, which a well-understood parameter. For (fringe-balanced) dangling-min BSTs, such a simple correspondence does not hold.  $S_n$  is an essential parameter for the linear-search part of SPINESEARCH.

$L_n$  is the number of keys in the leftmost leaf; by definition we have the following simple fact.

**Remark C.1**  $0 \leq \mathbb{E}[L_n] \leq w - 1 = O(1)$ . □

Note that by symmetry,  $\mathbb{E}[L_n]$  is actually the average number of keys in *any* leaf; a parameter that we will meet again when analyzing the memory usage of jumplists.

Note that distribution of  $P_n$  has a subtle complication, namely that even conditional on  $\mathbf{J}$ , the quantities  $S_n$ ,  $S_{J_1}$  and  $P_{J_1}$  are *not* independent: all consider the same left subtree! Our recurrence above does not reflect this; we actually have instead of  $P_{J_1}$  the path lengths of the left subtree *conditional* on having spine length  $S_{J_1}$  and leftmost leaf

size  $L_n$ ; similar dependencies exist for  $P_{J_2}$ . We refrained from doing so to avoid cluttering the equation; we will only compute the expected value here anyway. There, thanks to linearity, these dependencies can be ignored.

Using the distributional master theorem (DMT) (Theorem C.2) an asymptotic solutions for the expected values is quite easy to determine.

**Lemma C.4:**  $\mathbb{E}[S_n] \sim \frac{1}{2(H_{k+1} - H_{t+1})} \ln n$ .

**Proof:** We use Theorem C.2 with  $C_n \equiv S_n$ ,  $s \equiv 1$ ,  $n_0 \equiv w$ ,  $T_n \equiv 1$ , (i.e.,  $K \equiv 1$ ,  $\alpha \equiv 0$ ,  $\beta \equiv 0$ ),  $A_r \equiv 1$ , and with  $Z_1^* \equiv D_1$  the density convergence condition (20) is fulfilled, see Lemma C.3. We find  $H = 1 - \mathbb{E}[D_1^0] = 0$ , so Case 2 applies and yields the claim with  $\tilde{H} = -\sum_{r=1}^s \mathbb{E}[\ln(D_r)] = 2(H_{k+1} - H_{t+1})$ ; for the latter expectation, see the proof of Proposition 2.54 in [16].  $\square$

**Theorem C.5:**

*The expected number of key comparisons of a spine search in a random median-of- $k$  jumplist with leaf size  $w$  is*

$$\frac{\mathbb{E}[P_n]}{n+1} \sim \frac{1}{H_{k+1} - H_{t+1}} \ln n, \quad (n \rightarrow \infty),$$

when each of the  $n+1$  gaps is equally likely to be requested, and  $k$  and  $w$  are fixed constants.

**Proof:** We use the distributional master theorem (DMT) with  $C_n \equiv P_n$ , i.e.,  $s = 2$  and  $T_n \equiv (n+1) + (S_n + L_n + 1) + S_{J_1}$ . By Lemma C.4, all but the first summand are in  $\mathbb{E}[T_n]$  are actually in  $O(\log n)$ , so from the initially complicated toll function, only  $\mathbb{E}[T_n] \sim n$  remains in the leading term as  $n \rightarrow \infty$ . We thus have  $K \equiv 1$ ,  $\alpha \equiv 1$ ,  $\beta \equiv 0$ , in the DMT. The coefficients  $A_r \equiv 1$  for  $r = 1, 2$ , and with  $Z_1^* = D_1$  (as before). We find again  $H = 0$ , so Case 2 applies and yields the claim with  $\tilde{H} = -\sum_{r=1}^s \mathbb{E}[D_r \ln(D_r)] = H_{k+1} - H_{t+1}$ ; for the latter expectation, see the proof of Proposition 2.54 in [16].  $\square$

### C.3. Successful Searches

In the last section, we analyzed the external path length of median-of- $k$  jumplists; this corresponds to the averaged costs of SPINESEARCH. Since SPINESEARCH actually returns the last node with a strictly smaller key than the one sought, we need one additional (equality) comparison between the sought key  $x$  and the key of the successor of that node. The cost of successful and unsuccessful searches are then identical.

If we assume three-way comparisons, i.e., one comparison tells us in one shot whether  $x < y$ ,  $x = y$  or  $x > y$  holds, one can usually stop searching earlier if the key is found. Certainly, the costs for SPINESEARCH as analyzed above are an upper bound for the costs of this improved strategy; in fact, one can show that the leading-term asymptotic of the the path length is not affected by this optimization. We therefore do not discuss it further.

#### C.4. Insert

Inserting an element into a jumplist consists of three steps; first we spine-search the element in the list. Unless we find the element to already be present, we then add the new element to the backbone at the position returned by the search. Finally, we have to restore the proper distribution of the jump pointer lengths (RESTOREAFTERINSERT).

We already analyzed the cost of searches, and the second step only changes a constant number of pointers, so only the third part has to be discussed. RESTOREAFTERINSERT essentially retraces the search path to the element, and conditional on random coin flips, rebuilds at most one subtree using REBALANCE. Since REBALANCE has to visit each node in this subtree, its running time will be linear in the size of the corresponding subtree. We will therefore simply count the number of *rebalanced elements*, which is given by the second parameter,  $m$ , passed to REBALANCE.

The steps taken by RESTOREAFTERINSERT depend on the position of the newly inserted element; we denote here by  $R$  the rank of the gap the new element is inserted into. When the current sublist has  $m$  nodes, we have  $R \in [0..m]$ . Similar as for searches, we consider the average costs of insertion when all possible gaps are equally likely to be requested.

Unlike for searches, however, the distribution of  $R'$  in subproblems is *not* uniform even if  $R$  is: a close inspection of RESTOREAFTERINSERT reveals that (a) recursive calls in the jump-sublist always have  $R' \geq 1$ , and (b)  $R = 0$  and  $R = 1$  yield  $R' = 0$  in the recursive call in the next-sublist; in fact, once  $R = 0$  holds, we get this rank in all later recursive calls. We can therefore handle this disturbance by splitting the cases  $R = 0$  and  $R \geq 1$ ;

Also note that for the topmost call to RESTOREAFTERINSERT,  $R = 0$  is not possible, since no insertion before the header with dummy-key  $-\infty$  is possible. This means that initially  $R \stackrel{\mathcal{D}}{=} \mathcal{U}[1..m]$  holds. Recall that a jumplist on  $m$  nodes stores only  $n = m - 1$  keys, so that there are only  $n + 1 = m$  possible gaps initially. We obtain the following distributional recurrence for  $B_m^{\text{ins}}$ , the random number of rebalanced elements during insertion into the  $R$ th gap in a randomized median-of- $k$  jumplist with  $m$  nodes. (Note that unlike in the pseudocode,  $m$  is here the number of nodes in the jumplist *before* the insertion.)

$$B_m^{\text{ins}} \stackrel{\mathcal{D}}{=} F \cdot (m + 1) + (1 - F) \left( \mathbb{1}_{\{R=1\}} B_{J_1}^{\text{ins}0} + \mathbb{1}_{\{2 \leq R \leq J_1+1\}} B_{J_1}^{\text{ins}} + \mathbb{1}_{\{R \geq J_1+2\}} B_{J_2}^{\text{ins}} \right), \quad (m > w), \quad (31)$$

$$\text{where } R \stackrel{\mathcal{D}}{=} \mathcal{U}[1..m], \quad F \stackrel{\mathcal{D}}{=} B\left(\frac{k}{m-1}\right), \quad (32)$$

$$B_m^{\text{ins}} \stackrel{\mathcal{D}}{=} B_m^{\text{ins}0} \stackrel{\mathcal{D}}{=} [m = w] \cdot (m + 1), \quad (m \leq w), \quad (33)$$

$$B_m^{\text{ins}0} \stackrel{\mathcal{D}}{=} F \cdot (m + 1) + (1 - F) \cdot B_{J_1}^{\text{ins}0}, \quad (m > w), \quad (34)$$

$$\text{where } F \stackrel{\mathcal{D}}{=} B\left(\frac{k}{n-1}\right), \quad (35)$$

$$\mathbf{J} = (J_1, J_2) = \mathbf{I} + \mathbf{t} + 1, \quad (36)$$

$$\mathbf{I} = (I_1, I_2) \stackrel{\mathcal{D}}{=} \text{DirMult}(n - 2 - k; t + 1, t + 1). \quad (37)$$

In all cases,  $B_m$  terms on the right-hand side denote independent copies of the family of random variables and  $R$  and  $F$  are independent of  $B_n$  and  $\mathbf{J}$ .

**Lemma C.6:**  $\mathbb{E}[B_m^{\text{ins}0}] \sim \frac{k}{H_{k+1} - H_{t+1}} \ln m.$

**Proof:** We use once more the distributional master theorem. As before,  $Z_1^* \equiv D_1$  and the density convergence condition is satisfied (Lemma C.3). We have  $\mathbb{E}[T_n] \equiv \mathbb{E}[F(n+1)] \sim k$  and  $\mathbb{E}[A_1^{(n)}] \equiv \mathbb{E}[1 - F] = 1 \pm O(n^{-1})$ . Since  $H = 0$ , Case 2 applies and with  $\tilde{H} = -\mathbb{E}[\ln D_1] = H_{k+1} - H_{t+1}$  follows the claim.  $\square$

**Theorem C.7:**

*The expected number of rebalanced elements during the insertion of an element into a random median-of- $k$  jumplist with leaf size  $w$  over  $n$  keys is*

$$\mathbb{E}[B_{n+1}^{\text{ins}}] \sim \frac{k}{H_{k+1} - H_{t+1}} \ln n, \quad (n \rightarrow \infty),$$

when each of the  $n + 1$  possible gaps to insert a new element is equally likely, and  $k$  and  $w$  are fixed constants.

**Proof:** Towards applying the DMT on  $C_n \equiv B_n^{\text{ins}}$ , we compute

$$\mathbb{E}[T_n] \equiv \mathbb{E}[F(n+1) + (1-F)\mathbb{1}_{\{R=1\}}B_n^{\text{ins}0}] \quad (38)$$

$$= \frac{k(n+1)}{n-1} + \frac{n-1-k}{n-1} \frac{1}{n} \mathbb{E}[B_n^{\text{ins}0}] \quad (39)$$

$$\stackrel{\text{Lemma C.6}}{=} k \pm O(n^{-1} \log n). \quad (40)$$

As usual, we have  $Z_r^* = D_r$  ( $r = 1, 2$ ) and the density convergence condition is fulfilled by Lemma C.3. For the coefficients of the recursive terms holds

$$\mathbb{E}\left[A_1^{(n)} \mid Z_r^{(n)} \in (z - \frac{1}{n}, n]\right] \equiv \mathbb{P}[2 \leq R \leq J_1 + 1 \mid Z_r^{(n)} \in (z - \frac{1}{n}, n)] \quad (41)$$

$$= \frac{J_1}{n} = Z_1^{(n)} \quad (42)$$

$$= z \pm O(n^{-1}) \quad (43)$$

and similarly

$$\mathbb{E}[A_2^{(n)} \mid Z_r^{(n)} \in (z - \frac{1}{n}, n)] \equiv \mathbb{P}[R \geq J_1 + 2 \mid \mathbf{D}] \quad (44)$$

$$= z \pm O(n^{-1}) \quad (45)$$

so that  $a_1(z) = a_2(z) = z$ . Since  $H = 1 - \sum_{r=1}^2 \mathbb{E}[D_r] = 0$ , we again have Case 2 and find  $\tilde{H} = -\sum_{r=1}^2 \mathbb{E}[D_r \ln D_r] = H_{k+1} - H_{t+1}$ . This proves the claim.  $\square$

## C.5. Delete

Like insertion, deleting an element from a jumplist consists of three stages: first we search the element, and, unless it is not present, we then remove it from the list and finally restore

the distribution in the resulting jumplist (RESTOREAFTERDELETE). As for insertion, we analyze the size of the sublist  $B_m^{\text{del}}$  that is rebuilt using REBALANCE in this last step when the rank of the deleted element is chosen uniformly; as for insertion, we initially have  $R \stackrel{\mathcal{D}}{=} \mathcal{U}[2..m]$  since the dummy key  $-\infty$  in the header cannot be deleted. In recursive calls, also  $R = 1$  is possible, and we remain in this case for good whenever we enter it once. We can thus characterize the deletion costs using the two quantities  $B_m^{\text{del}}$  and  $B_m^{\text{del1}}$ . As for insertion,  $m$  is the “old” size of the jumplist, i.e., the number of nodes *before* the deletion.

$$B_m^{\text{del}} \stackrel{\mathcal{D}}{=} F \cdot (m-1) + (1-F) \left( \mathbb{1}_{\{R=2\}} B_{J_1}^{\text{del1}} + \mathbb{1}_{\{3 \leq R \leq J_1+1\}} B_{J_1}^{\text{del}} + \mathbb{1}_{\{R \geq J_1+3\}} B_{J_2}^{\text{del}} \right), \quad (m > w) \quad (46)$$

$$\begin{aligned} \text{where } R &\stackrel{\mathcal{D}}{=} \mathcal{U}[2..m] \\ \text{and (conditional on } \mathbf{J}) \quad F &\stackrel{\mathcal{D}}{=} \begin{cases} B\left(\frac{t}{J_1-1}\right), & R \leq J_1+1; \\ 1, & R = J_1+2; \\ B\left(\frac{t}{J_2-1}\right), & R \geq J_1+3, \end{cases} \end{aligned} \quad (47)$$

$$B_m^{\text{del}} \stackrel{\mathcal{D}}{=} B_m^{\text{del1}} \stackrel{\mathcal{D}}{=} [m=w] \cdot 1, \quad (m \leq w), \quad (48)$$

$$B_m^{\text{del1}} \stackrel{\mathcal{D}}{=} F \cdot (m-1) + (1-F) \cdot B_{J_1}^{\text{del1}} \quad (49)$$

$$\text{where (conditional on } \mathbf{J}) \quad F \stackrel{\mathcal{D}}{=} B\left(\frac{t}{J_1-1}\right), \quad (50)$$

$$\mathbf{J} = (J_1, J_2) = \mathbf{I} + \mathbf{t} + 1, \quad (51)$$

$$\mathbf{I} = (I_1, I_2) \stackrel{\mathcal{D}}{=} \text{DirMult}(m-2-k; t+1, t+1). \quad (52)$$

The (asymptotic) solution of these recurrences works very similar to the case of insertion. We first consider deleting leftmost keys

**Lemma C.8:** *If  $t \geq 1$ ,  $\mathbb{E}[B_m^{\text{del1}}] \sim \frac{k}{H_{k+1} - H_{t+1}} \ln m$ ; for  $t = 0$  we have  $\mathbb{E}[B_m^{\text{del1}}] \leq 1$ .*

**Proof:** Note that for  $t = 0$ , we have  $F = 0$  (almost surely) in each iteration, so the recurrence immediately collapses to its initial condition, which is at most 1. In the following, we now consider  $t \geq 1$ .

The proof will ultimately use the DMT (Theorem C.2) on  $C_n \equiv B_n^{\text{del1}}$ , but we need a few preliminary results to compute the toll function  $\mathbb{E}[T_n] \equiv \mathbb{E}[F(n-1)]$ . First, we note the following truncated development of  $\frac{c}{x+c}$  in terms of (negative) falling factorial powers of  $x$ :

$$\forall t \in \mathbb{N}_{\geq 1} \quad \forall n \geq 0 \quad : \quad \frac{t}{n+t} = tn^{-1} \pm t(t-1)n^{-2} \quad (53)$$

(The proof is elementary.)

We use the shorthand  $a = b \pm d$  to mean  $b - d \leq a \leq b + d$  here and throughout. Now, we compute the expectation of  $F$  conditional on  $\mathbf{I} = \mathbf{J} - \mathbf{t} - 1$ .

$$\mathbb{E}[F \mid \mathbf{I}] = \frac{t}{J_1 - 1} \quad (54)$$

$$= \frac{t}{I_1 + t} \quad (55)$$

$$\stackrel{(53)}{=} t \cdot I_1^{-1} \pm t(t-1) \cdot I_1^{-2} \quad (56)$$

Next, we use the stochastic representation

$$\text{BetaBin}(n, \alpha, \beta) \stackrel{\mathcal{D}}{=} \text{Bin}(n, D_1) \quad \text{conditional on } D_1 \stackrel{\mathcal{D}}{=} \text{Beta}(\alpha, \beta). \quad (57)$$

Taking expectations over  $I_1 \stackrel{\mathcal{D}}{=} \text{Bin}(\eta, D_1)$  with  $\eta = m - 2 - k$ , but still keeping  $\mathbf{D} = (D_1, D_2) = (D_1, 1 - D_1)$  fixed and using Lemma C.1, we find

$$\mathbb{E}[F \mid \mathbf{D}] = \frac{t}{\eta + 1} D_1^{-1} (1 - D_2^{\eta+1}) \pm t(t-1) D_1^{-2} \eta^{-2} \quad (58)$$

Finally, we also compute the expectation w.r.t.  $\mathbf{D} \stackrel{\mathcal{D}}{=} \text{Dir}(t+1, t+1)$ ; note that for  $t \geq 2$ ,  $\mathbb{E}[D_1^{-2}]$  exists and has a finite value independent of  $n$ ; whereas for  $t = 1$ , the error term is actually exactly zero. We hence find in both cases using [16, Lemma 2.30], the “powers-to-parameters” rule:

$$\mathbb{E}[F] = \frac{t}{\eta + 1} \mathbb{E}[D_1^{-1}] - \frac{t}{\eta + 1} \mathbb{E}[D_1^{-1} D_2^{\eta+1}] \pm O(\eta^{-2}) \quad (59)$$

$$= \frac{t}{\eta + 1} \frac{k}{t} - \frac{t}{\eta + 1} \frac{(t+1)^{\overline{\eta+1}}}{t(k+1)^{\overline{\eta}}} \pm O(\eta^{-2}) \quad (60)$$

$$= \frac{k}{\eta + 1} - \frac{(t+1)(t+2)}{(\eta+1)(\eta+2)} \underbrace{\frac{(t+3)^{\overline{\eta-1}}}{(k+2)^{\overline{\eta-1}}}}_{<1} \pm O(\eta^{-2}) \quad (61)$$

$$= \frac{k}{\eta + 1} \pm O(\eta^{-2}) \quad (62)$$

With this we finally get  $\mathbb{E}[T_n] \equiv \mathbb{E}[F(n-1)] = k \pm O(n^{-1})$ . We are now in the position to apply the DMT with  $\mathbb{E}[A_1^{(n)} \mid Z_1^{(n)} \in (z - \frac{1}{n}, z)] = 1 \pm O(n^{-1})$  and  $Z_1^* = D_1$ . We find  $H = 0$  and obtain the claim by Case 2 with  $\tilde{H} = -\mathbb{E}[\ln D_1] = H_{k+1} - H_{t+1}$ .  $\square$

**Theorem C.9:**

*The expected number of rebalanced elements during the deletion of an element from a random median-of- $k$  jumplist with leaf size  $w$  over  $n$  keys is*

$$\mathbb{E}[B_{n+1}^{\text{del}}] \sim \frac{k}{H_{k+1} - H_{t+1}} \ln n, \quad (n \rightarrow \infty),$$

*when each of the  $n$  keys is equally likely to be deleted, and  $k$  and  $w$  are fixed constants.*

**Proof:** We start with computing the conditional expectation of  $F$ , the coin flip indicator.

$$\mathbb{E}[F \mid \mathbf{J}] = \frac{J_1}{n-1} \cdot \frac{t}{J_1-1} + \frac{1}{n-1} \cdot 1 + \frac{J_2-1}{n-1} \cdot \frac{t}{J_2-1} \quad (63)$$

$$= \frac{2t+1}{n-1} + \frac{1}{n-1} \cdot \frac{t}{J_1-1} \quad (64)$$

$$\rightsquigarrow \mathbb{E}[F] \stackrel{(62)}{=} \frac{2t+1}{n-1} + \frac{1}{n-1} \cdot \frac{k}{\eta+1} \pm O(n^{-3}) \quad (65)$$

$$= \frac{k}{n-1} \pm O(n^{-2}) \quad (66)$$

Towards applying the DMT on  $C_n \equiv B_n^{\text{del}}$ , we compute

$$\mathbb{E}[T_n] \equiv \mathbb{E}[F(n-1) + (1-F)\mathbf{1}_{\{R=2\}}B_n^{\text{del}}] \quad (67)$$

$$\stackrel{\text{Lemma C.8}}{=} k \pm O(n^{-1} \log n) \quad (68)$$

As usual, we have  $Z_r^* = D_r$  and the density convergence condition is fulfilled. For the coefficients of the recursive terms holds similarly as in Theorem C.7  $a_1(z) = a_2(z) = z$ . Once more we have  $H = 0$  and Case 2 applies. With  $\tilde{H} = -\sum_{r=1}^2 \mathbb{E}[D_r \ln D_r] = H_{k+1} - H_{t+1}$ , the claim follows.  $\square$

## C.6. Memory Requirements

To estimate the amount of memory used by a typical implementation to store a jumplist on  $n$  keys, we make the following assumptions. A pointer requires one word of storage, and so does an integer that can take values in  $[0..n+1]$ . We do not count memory to store the keys since any (general-purpose) data structure has to store them. This means that a plain node requires 1 word of (additional) storage, whereas a jump node needs 3 words. If  $A_n$  denotes the (random) number of jump nodes, excluding the dummy header, of a random median-of- $k$  jumplist with leaf size  $w$ , its memory requirement (on top of the  $n$  keys) is simply  $3(A_n + 1) + 1(n - A_n)$ .

It remains to determine  $A_n$ . Each of the  $A_n$  jump nodes corresponds to exactly one internal node in the corresponding fringe-balanced dangling-min BST over  $n$  keys. A distributional recurrence is easy to set up:

$$A_n \stackrel{\mathcal{D}}{=} 1 + A_{J_1} + A_{J_2}, \quad (n > w-1), \quad (69)$$

$$A_n \stackrel{\mathcal{D}}{=} 0, \quad (n \leq w-1) \quad (70)$$

$$\mathbf{J} = (J_1, J_2) = \mathbf{I} + \mathbf{t}, \quad (71)$$

$$\mathbf{I} = (I_1, I_2) \stackrel{\mathcal{D}}{=} \text{DirMult}(n-1-k; t+1, t+1), \quad (72)$$

For  $A_n$ , the DMT only gives us  $\mathbb{E}[A_n] = O(n)$ , since we end up in Case 3. It is easy to see that  $\mathbb{E}[A_n]$  is at least linear for constant  $w$ , but a precise leading-term asymptotic seems indeed very hard to obtain in this case.<sup>3</sup> The recurrence for  $A_n$  is, however, very similar

<sup>3</sup> Readers familiar with generating functions and the Euler differential equation for them that we obtain for the very similar Quicksort recurrence may find it interesting that in this case, the differential equation for the generating function is *not* Euler equation. Hence there is little hope to obtain a solution for the generating function for  $t \geq 1$ .

to that for the number of partitioning steps in median-of- $k$  Quicksort with Insertionsort threshold  $w - 1$ —or equivalently, the number of internal nodes in median-of- $k$  BSTs with leaf buffer  $w - 1$ . The only difference is that one there has  $\mathbf{I} \stackrel{\mathcal{D}}{=} \text{DirMult}(n - k; t + 1, t + 1)$ :  $n - k$  instead of  $n - 1 - k$ . By monotonicity,  $\mathbb{E}[A_n]$  is clearly at most as the number of partitioning steps in Quicksort since also the subproblems sizes are smaller. We therefore obtain the following statement.

**Theorem C.10:**

*The asymptotic fraction of jump nodes in a random median-of- $k$  jumplist with leaf size  $w$ , where  $k$  and  $w$  are fixed constants, is at most*

$$\frac{1}{(w + 1)(H_{k+1} - H_{t+1})}.$$

*The average number of additional words per key in such a jumplist thus is*

$$1 + \frac{2}{(w + 1)(H_{k+1} - H_{t+1})}.$$

**Proof:** The number of partitioning steps in median-of- $k$  Quicksort with Insertionsort threshold  $M$  is  $1/((M + 2)(H_{k+1} - H_{t+1}))n \pm O(1)$ , see, e.g., Hennequin [5, p. 327]. Setting  $M = w - 1$  and our discussion above yield the claim.  $\square$